



# Supermicro Total Solution for Lustre on OpenZFS

High performance file system based on open software defined storage accelerates innovation and redefines storage economics

## Scale-Out Storage - Lustre

Lustre is the Parallel file system of choice for High Performance Computing (HPC) and large file applications. To meet the capacity and throughput requirements of HPC workloads, Lustre has traditionally required adoption of custom proprietary storage products leading to vendor lock-in and reduced innovation. The evolution of software defined high availability platforms like OpenZFS running on Linux now make it possible to meet the performance requirements of HPC on open industry standard x86 platforms, reducing storage costs by up to 90%. Supermicro in Partnership with Intel® and some of the worlds leading HPC integrators offers a total solution for Lustre on OpenZFS with Supermicro's industry leading hardware, software and services infrastructure.

## Open Industry Standard <Software/Defined> Architecture

- Tested and Validated Lustre Open solution based on workload optimized Supermicro Systems
- Reduce storage costs by up to 90% with Lustre on OpenZFS based on Supermicro open industry standard hardware and software
- Standard Configuration delivers industry leading storage density (90 Bay 4U) with fully redundant High Availability design
- Performance Configuration delivers unparalleled All-Flash NVMe performance in a 2U Design with full system (dual node) and data path (dual-path NVMe) redundancy

## Best in Class System and Software Infrastructure

- Leverages the best of Supermicro's expansive Server and Storage product portfolio
- Based on hardened Intel Enterprise Edition for Lustre\* Software on OpenZFS optimized for Supermicro systems increasing single node performance by 30%
- Flexible architecture providing a high performance, hot swappable and no single point of failure design
- High bandwidth utilizing 12Gb/s SAS3 and 56Gb/s FDR InfiniBand

## Total Solution (Software, System and Service)

Pre-designed and pre-tested Total Solution for Lustre on OpenZFS provides customers the insight on where and how to deploy Lustre on OpenZFS reducing complexity and risk, controlling costs and accelerating solution time to market.

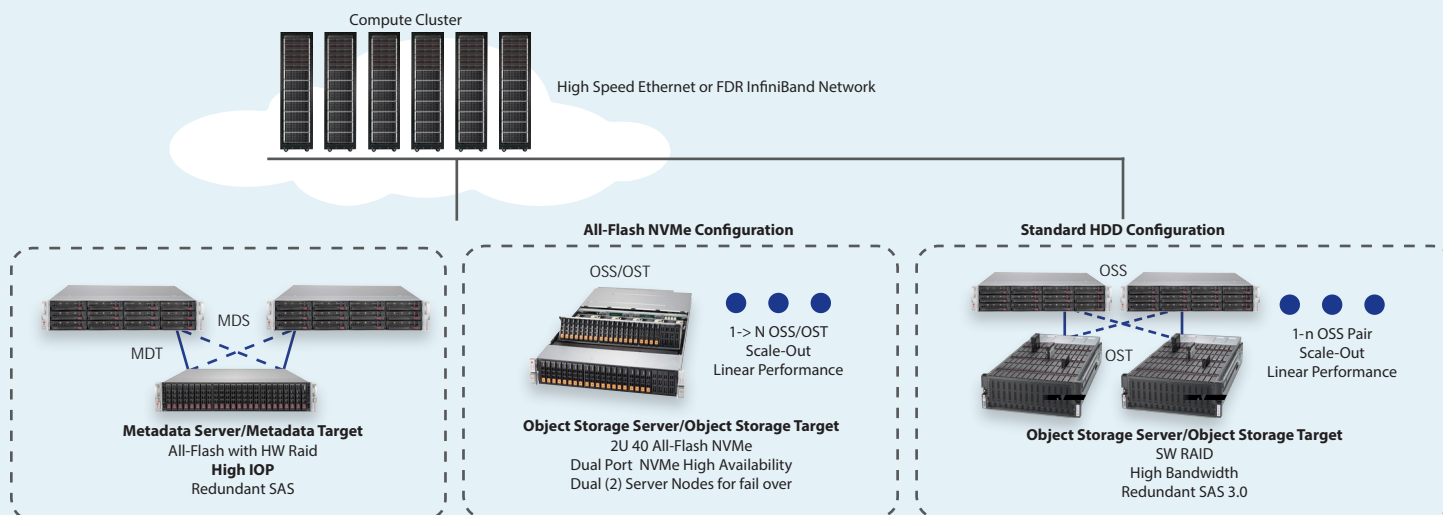
## LUSTRE TOTAL SOLUTION READY HARDWARE



2U 40 Bay Dual Node Server with  
Dual Port All-Flash NVMe




4U 3.5" 90 Bay Dual Expander JBOD





# Total Solution for Lustre Configurations



**Metadata and Mgmt Server Pod (MDS/MGS)**

	<b>Object Storage Server All-Flash NVMe</b>	<b>Object Storage Server Pod HDD</b>	<b>Metadata and Mgmt Server Pod (MDS/MGS)</b>
<b>Description</b>	<ul style="list-style-type: none"><li>• <b>OSS:</b> Multiple systems can be added to meet capacity requirements</li><li>• Linear scaling of throughput and capacity</li></ul>	<ul style="list-style-type: none"><li>• Multiple Pods can be added to meet capacity requirements</li><li>• Linear scaling of throughput and capacity</li></ul>	<ul style="list-style-type: none"><li>• Single MDS/MGS Pod can support hundreds of OSS Pods and thousands of clients</li></ul>
<b>Capacity/Pod</b>	30 TB raw capacity using 750GB media	720TB raw capacity using 8TB media	3.2TB raw capacity, user upgradable
<b>Form Factor</b>	<ul style="list-style-type: none"><li>• 2U Configuration</li><li>• Redundant Server Nodes</li><li>• Dual-Path Redundant NVMe</li><li>• 40-bay Storage Target (OST)</li></ul>	<ul style="list-style-type: none"><li>• 8U Configuration</li><li>• Redundant Object Storage Servers</li><li>• 1x 90-bay Object Storage Target (OST) (16x lanes per node)</li></ul>	<ul style="list-style-type: none"><li>• 6U Configuration</li><li>• Redundant Metadata servers</li><li>• shared SAS3 JBOD (8x lanes per node)</li></ul>

## Technical Specifications

### MDS/MGS

<b>CPU/Memory</b>	Dual E5-2667 v4, 3.2 GHz, 8 core/ 128GB
<b>Fabric</b>	User definable supporting Ethernet, InfiniBand and Omni-Path Interconnects up to 100Gb/s
<b>Storage Interconnect</b>	12Gb/s (SAS 3) LSI/Avago Syncro hardware RAID controllers
<b>Metadata Target</b>	2U 24 2.5" Drive JBOD with shared redundant expander (8X lanes/node)

### OSS/OST (HDD)

<b>CPU/Memory</b>	Dual E5-2667v4, 3.2 GHz, 8 core/ 128GB
<b>Fabric</b>	User definable supporting Ethernet, InfiniBand and Omni-Path Interconnects up to 100Gb/s
<b>Storage Interconnect</b>	12Gb/s (SAS 3) redundant IT mode controllers
<b>Object Data Target</b>	4U 90 3.5" drive JBOD(s) with shared SAS 3 redundant expander (16x lanes/node)

### OSS/OST

<b>Nodes</b>	2 Redundant Server Nodes in 2U
<b>CPU/Memory</b>	Dual E5-2667v4, 3.2 GHz, 8 core/ 128GB
<b>Fabric</b>	User definable supporting Ethernet, InfiniBand and Omni-Path Interconnects up to 100Gb/s
<b>Storage Interconnect</b>	Redundant Switched PCI-E 3.0 Architecture
<b>Drive Support</b>	2U x 40 All-Flash NVMe

## Management Software

- Intel Manager for Lustre
- Open Industry Standard server management (SSM)
- Remotely manage hardware and OS configuration, health & power consumption of nodes in cluster

## Support & Integration

- Flexible and customizable service level agreements (SLA): 4-Hour and Next Business Day Onsite Service
- Consulting service on solution architecture design
- Engineering support on installation, configuration and testing

For more Information, call your Supermicro sales representative.

[www.Supermicro.com/solutions](http://www.Supermicro.com/solutions)